Identifying the Evolutionary Progression of Color from Crosslinguistic Data

Julia Watson
Department of Computer Science
University of Toronto
(jwatson@cs.toronto.edu)

Barend Beekhuizen
Department of Language Studies
University of Toronto
(barend.beekhuizen@utoronto.ca)

Suzanne Stevenson
Department of Computer Science
University of Toronto
(suzanne@cs.toronto.edu)

Abstract

We present a novel statistical analysis of color categorization using a standard method from semantic typology. Our approach shows that crosslinguistic color naming data exhibits latent dimensions whose order of relative importance matches the evolutionary ordering of emergence of those distinctions. Moreover, we show that the importance ordering of these dimensions holds even when controlling for frequency of the distinctions by looking at languages within each stage of evolution. Additionally, we find that the extreme points of the latent color dimensions correspond well to a small set of "universal" focal colors. Thus we show that a simple mathematical method simultaneously derives a consistent match both to the evolutionary stages and to the universal foci.

Keywords: semantic universals; color naming; color evolution.

Introduction

Much work in cognitive science seeks to uncover the basis of human categorization of the world. Semantic typology in particular aims to discover crosslinguistic constraints and tendencies in the ways that lexical semantic systems parcel concepts into named categories. Research across a number of diverse domains – from color to spatial relations to cutting and breaking events (e.g., Berlin & Kay, 1969; Levinson et al., 2003; Majid et al., 2008) – have revealed seemingly universal dimensions that underlie the organization of such lexical categories. For example, there is substantial evidence that color lexicons are organized around a universal set of basic color categories, whose best exemplars – *focal colors*, or *foci* – are clustered within small areas of the perceptual color space (e.g., Berlin & Kay, 1969; Regier et al., 2005; though see, e.g., Roberson et al., 2000, for an alternative view).

The domain of color has been particularly fruitful in revealing such crosslinguistic commonalities. Indeed, research on color is unusual (if not unique) in semantic typology in having revealed another kind of universal as well – that of evolutionary stages of a domain-specific lexicon. Berlin & Kay (1969) proposed that, as the number of basic color terms increase in a language, the named color distinctions emerge in one of a small set of constrained orders; for example, separate terms for yellow and red appear in a language before green is split off from blue. This line of work has been extended to cover a broad range of data from many languages, and the specific proposal refined and adapted. While some counterexamples have been identified, and it has been recognized that some languages do not exhaustively partition the color spectrum, most extant languages largely follow a successive partitioning of the color space according to universal principles (see Kay et al., 2009, for a review).

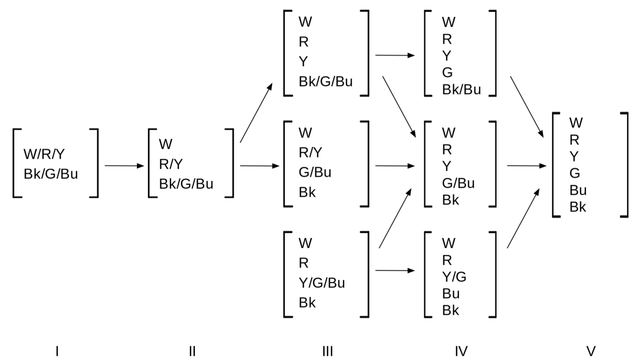

Figure 1: Evolutionary chart from Kay et al. (2009). W=white, R=red, Y=yellow, Bk=black, G=green, Bu=Blue.

Most evolutionary proposals focus on a core set of basic color categories, corresponding to the English terms *white*, *red*, *yellow*, *black*, *green*, and *blue*, because these 'primary colors' follow a consistent evolutionary path (Kay et al., 2009). Fig. 1 illustrates an influential proposal regarding the evolutionary sequence of languages, which we follow here. This diagram says that languages with only two color terms (Stage I) allocate those to the distinction between warm colors (White/Red/Yellow) and cool colors (Black/Green/Blue), while languages with three colors (Stage II) further distinguish White from Red/Yellow. Languages at Stages III through V can follow multiple pathways, depending on the further splits of Red, Yellow, Black, Green, and Blue.

Although modeling of crosslinguistic color data has revealed evidence of the various stages in Fig. 1 (e.g., Regier et al., 2007; Lindsey & Brown, 2009; Jäger, 2012), to our knowledge such work has not (yet) shown how the *ordering* of such stages could be derived from synchronic color naming data alone. Moreover, work on the evolutionary stages has typically focused on the partitioning of the space, and has not linked those stages to the nature and role of focal colors. That is, we know of no single model or method of analysis over the crosslinguistic color data that has derived both a consistent match to the evolutionary stages and to the universal focal areas, showing if or how these two concepts are linked.

Here we show that a standard analysis method from semantic typology, which has been used in work on color as well as

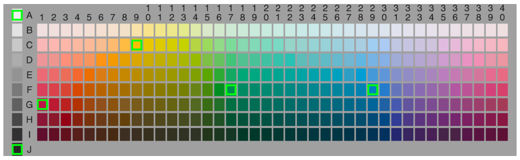

Figure 2: Munsell chart; universal foci (Regier et al., 2005)

in various other semantic domains, can simultaneously derive both the ordering of the evolutionary stages and the universal focal colors, revealing the latter as the drivers of the evolutionary distinctions. To preview our results, we use a simple statistical method to identify the important dimensions of the semantic space of color as represented by the crosslinguistic data. We find that the extremes of the identified dimensions correspond to focal areas of color, and, crucially, that the *importance ordering* of the dimensions corresponds to the *evolutionary ordering* of color distinctions. In addition, we perform the first analysis of subsets of languages at different evolutionary stages, and show that this importance ordering is not simply a by-product of the frequency of color distinctions in a mixed set of languages. We thus provide quantitative confirmation that synchronic color naming patterns reveal an underlying semantic space whose dimensional salience mirrors the evolutionary stages, with "anchors" in focal colors.

Color Data, Foci, and Evolutionary Stages

The World Color Survey (WCS; Kay et al., 2009, http://www.icsi.berkeley.edu/wcs/data.html) is a rich data set, along with a comprehensive qualitative analysis, that underlies much crosslinguistic analysis of color categories. Speakers of 110 diverse languages were surveyed, and languages were manually assigned to an evolutionary stage or transition between stages (cf. Fig. 1). The WCS contains two types of data. First, *naming data* was gathered by asking speakers to provide a single color term for each of the 330 color chips in the Munsell chart (see Fig. 2). Second, *focal data* was collected by asking speakers to select the most representative example, or focal color, of each term.

Both the focal and naming data of the WCS have played a prominent role in semantic typological analyses of color, which seek to derive semantic universals from crosslinguistic usage data. In particular, a body of work has attempted to go beyond qualitative analyses to provide precise mathematical underpinnings for such universals. Our work is in this vein, and we review related research below. Other quantitative analyses have attempted to link the semantic universals apparent from the WCS data to perceptual and/or communicative aspects of cognition. This is not the goal of our work here, but we refer to such research where relevant.

It is a striking finding that the distributions of the focal colors across all of the WCS languages cluster in small areas of the color space, corresponding to the six basic English focal colors, *white*, *red*, *yellow*, *black*, *green*, and *blue* (MacLaury, 1997; Regier et al., 2005; Lindsey & Brown, 2006, see Fig. 2). Some claim that these "universal" focal colors are cognitively privileged areas of the mental representation of color (Heider, 1972; Regier et al., 2005), which play a crucial role in the evolution of color systems (Berlin & Kay, 1969); others propose that they are only epiphenomena of the desired placement of color category boundaries (Roberson et al., 2000).

In their perceptual account, Jameson & D'Andrade (1997) suggest a middle ground in which the focal colors arise due to the nature of categorization in an irregular perceptual space. Abbott et al. (2012) operationalize this approach using Bayesian inference over perceptual color categories whose extents are determined by the WCS naming data. They find a good match between the representative members of these named color categories and the WCS focal data. This suggests that foci may be derivative from color categories whose optimal boundaries are driven by universal properties of the perceptual space (e.g., Jameson & D'Andrade, 1997; Regier et al., 2007). However, the relation of such foci to the evolutionary distinctions among colors is not clear.

The WCS data has also been explored as a source of insight into the evolutionary stages of color term systems, as exemplified in Fig. 1. Lindsey & Brown (2006, 2009) apply clustering techniques over naming patterns to reveal universal constraints over color categories, as well as color naming "motifs" (ways of partitioning the color space), some of which correspond to stages in the evolutionary hierarchy. Jäger (2012) takes a complex, multi-step approach to applying PCA to the WCS naming data, after transforming it in various ways. He derives partitionings of the six basic colors, many (but not all) of which match those of the evolutionary stages. While these approaches use quantitative analyses of the WCS to derive aspects of the evolutionary partitions, none of them derives an *ordering* over the partitions. (Indeed, Lindsey & Brown (2006) explicitly note that their work should not be interpreted as evidence of evolutionary sequencing from synchronic data.)

By contrast, Zaslavsky et al. (2018) combine WCS naming data with a perceptual semantics to derive an order over the emergence of color categories. They assume that color categories are created to optimally balance lexical accuracy with the size of the lexicon. As more color categories are added, their emergence roughly reflects the ordering of categories in color evolution. However, the reliance on perceptual salience leads to some mismatches with the evolutionary stages (overestimating the prominence of yellow), and the method does not address the role of focal colors in the ordering.

The wealth of research analyzing the WCS motivates our exploration of whether this rich synchronic color naming data can directly reveal patterns of evolutionary development, and shed light on the role of focal colors in those stages. We aim for a mathematical method of analysis that is simple and straightforward, with the intention that such an approach would be readily applicable to other semantic domains.

Our Approach

The approach we take in this work complements and seeks to fill in some of the gaps noted in the above body of research. Our goal is to derive the evolutionary sequence from WCS data using a simple mathematical method – Principal Component Analysis (PCA) – that (along with other dimensionality reduction techniques) has been widely deployed in semantic typology, in diverse semantic domains including color (e.g., Majid et al., 2008; Jäger, 2012; Beekhuizen et al., 2014; Beekhuizen & Stevenson, 2018).

The novelty of our approach is two-fold. First, we extract latent dimensions of the WCS data *in order of importance*, yielding the first quantitative evidence of the evolutionary progression of color naming from the synchronic data. Second, we propose a natural interpretation of the "extremes" of the extracted dimensions as focal areas of color, which indeed show a strong match to empirical foci. Thus, we achieve a simultaneous match of the evolutionary ordering and the focal colors, which has not been shown before. Moreover, we do so with a very simple and straightforward use of PCA, in contrast to other methods that require much more involved mathematical transformations of the data (as in Jäger, 2012).

Our motivation is as follows. If most languages have followed a consistent and small set of orderings in the diachronic emergence of colors, those orderings must be determined by the relative importance of various perceptual/behavioral/cultural/communicative influences (e.g., Jameson & D'Andrade, 1997; Kay et al., 2009; Gibson et al., 2017; Holmes & Regier, 2017; Zaslavsky et al., 2018). Regardless of the source of these influences, if they play a role in the evolution of color systems, they may impact synchronic use of color terminology, and with the same relative importance. Note that this is not necessarily the case; for example, just because a language at Stage V has gone through Stages I through IV does not mean that the current color naming patterns of that language will reflect that a distinction made in an earlier stage (e.g., of Red vs. Yellow) is more important than a final distinction made in Stage V (e.g., of Green vs. Blue). That is, it is an open question whether the factors that exert evolutionary pressure to *create* new terms play a role in how terms are *deployed* in synchronic naming.

Experimental studies suggest that it may indeed be the case that evolutionary factors play a role in cognitive processing of colors by individuals. For example, Holmes & Regier (2017) found that English speakers show a categorical perception effect for the warm–cool distinction of Stage I, even though "warm" and "cool" are not basic color terms in English. Moreover, when English speakers group colors into K categories, the divisions they make roughly follow the evolutionary splits – i.e., with $K = 2$, they select a warm–cool separation, as in Stage I of evolution, with $K = 3$ they add a further distinction of white as in Stage II, etc. (Boster, 1986; Xu et al., 2013). Thus, English speakers are apparently sensitive to the evolutionary factors – and their relative ordering of importance – in color category processing.

Our goal here is to see whether actual color naming behavior, across the many diverse languages of the WCS, show this synchronic realization of the evolutionary influences. Specifically, given a suitable representation of the semantic space of synchronic color naming patterns, we use PCA over this data to extract dimensions of the data in order of importance, and examine whether those dimensions and their relative importance match the evolutionary stages proposed in the literature.

As a suitable representation of the color naming data, we follow a straightforward and standard practice in semantic typology. Specifically, we create a color chip by color term matrix using the color naming data from the WCS (Beekhuizen & Stevenson, 2018). Intuitively, such a matrix forms a semantic space over color, where each row can be viewed as a vector representation of the meaning of a color chip, as determined by the aggregate naming patterns in the data.

Applying PCA to such a matrix finds the latent dimensions characterizing the semantics of color across languages. Moreover, we take advantage of the interpretability of PCA dimensions, which means that points with a minimum or maximum value for a dimension are the most "extreme" example of the property that that dimension captures. Such points represent the "corners" of the data in the space (cf. Fig. 3), which are an indication of the key distinction each dimension is enforcing. We can thus examine these extremes to see if they correspond to the focal colors that have been proposed to "anchor" color categories (Regier et al., 2005).

Methods

Data matrices and PCA. We first create a color chip by term matrix over the naming data. Each cell records the (normalized) number of speakers in a language that used that term for that chip. This matrix compiles the naming data from all or selected subsets of languages in the WCS (as noted below). Thus we create matrices with 330 rows (one per chip in the Munsell chart) and up to 2223 columns (the number of color terms across all WCS languages).

We apply PCA to the resulting matrices to extract the most important dimensions. PCA identifies dimensions in the order along which the data shows the most variance, so the amount of variance accounted for represents the importance of that dimension. As we are looking for important dimensions that could relate to evolutionary development, we only consider dimensions that account for at least 5% of the variance in the data. (In almost all cases, this corresponded to a natural dropping off point in the accounted-for variance.)

For the first three such dimensions, we can plot the data for visualization purposes; i.e., we can plot the 330 color chips as represented by the first dimension of the PCA, by the first two, or by the first three. As shown in Fig. 3, such plots reveal the structure in the data that the PCA finds.

Determining the extreme points. To better understand the dimensions extracted by the PCA, we want to identify the *extreme points* of each. Conceptually, these are the maximally distinguishable points in the data on that dimension; in our visualizations in Fig. 3, these correspond to the endpoints or "corners" of the plotted data. In Fig. 3, the extreme points in 1D correspond to the minimum and maximum values on the x axis; the extreme points in 2D correspond to the corners of a triangle; the extreme points in 3D are the top of the pyramid and the corners of its triangular base. As we will see in the results, such points generally correspond to focal colors.

We first collect the minimum and maximum points per dimension. However, because points can be tightly clustered at a "corner" of the space, we can end up with multiple extreme points when there is really only one "corner". For example, in the 2D plot in Fig. 3 (middle panel), the bottom right corner of the triangle is near the maximum for the x dimension *and* the minimum for the y dimension, so we might find two distinct points in that same small area. To avoid this, we consider all pairwise combinations of extreme points and merge those that are likely referring to the same "corners" of the space. Extreme points are considered to refer to the same "corner" if there is overlap in their $n = 15$ nearest neighbors, based on Euclidean distance. (We tried other values of n, which made little difference in the pattern of results.)

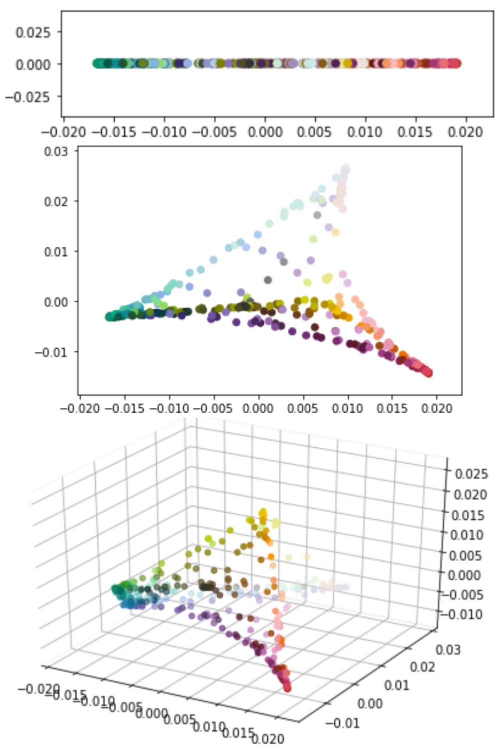

Figure 3: Plots of the 330 color chips in the 1D, 2D, and 3D PCA subspaces of the full WCS.

Visualizations of results. The results of our PCA analysis are a sequence of semantic subspaces defined by the extracted dimensions – the first (1D), the first+the second (2D), the first+the second+the third (3D), etc. – and the points that indicate the extreme "corners" of each subspace. To visually show the results, we plot the extreme points of each subspace (as triangles; or as diamonds for the merged points) in a Munsell chart, along with their closest neighbors (as circles whose size reflects their distance from the extreme point). These extreme point areas show the important color prototypes for each of the components of the PCA.

In addition, we visualize the extreme points as partitioning the PCA subspace, such that every color chip is allocated to the region of the space of its nearest extreme point. This yields a partitioned Munsell chart, with the number of partitions equal to the number of (merged) extreme points, or "corners", in the space. We label these partitions by the focal colors (green squares in the charts) occurring within them, whether they are extreme points or not. Thus each of the six focal colors is allocated to the region of its nearest extreme point, and we label a region by the focal colors it includes.

Fig. 4 shows examples of this visualization. For example, the White, Red, and Green extreme points shown in Fig. 4b for 2D correspond to the white, red, and green "corners" of the PCA plot for 2D shown in Fig. 3. The labels W, R/Y, and Bk/G/Bu correspond to the focal colors within each region as partitioned by the extreme points.

Analysis Over All WCS Languages

Using the methods above, our goal is to see whether a simple and straightforward application of PCA over the WCS naming data can simultaneously derive both the ordering of the evolutionary stages, as in Fig. 1, and the location of the universal foci, as in Fig. 2.

Results. We first apply our method over the full WCS color naming data. The first five extracted dimensions each account for more than 5% of the variance in the data; the results on the corresponding 1D–5D subspaces are in Fig. 4(a–e).

First, note that all of the primary extreme points for all dimensions of the PCA occur very close to the universal focal points of the 6 basic colors; see the triangles in Fig. 4. The extreme points corresponding to White, Black, and Green occur at the focus, Yellow and Blue adjacent to the focus, and Red two away from the focus. Fig. 4(f) shows the close match between our predicted focal colors and those of Abbott et al. (2012), who draw on both the color naming data and a representation of the named categories in perceptual space. It is important to emphasize that (as in Abbott et al., 2012) our results do not make use of the WCS focal color data, but only naming data. Given evidence for the universality of the focal colors as "anchors" for language-dependent naming of color regions, our results suggest that the extreme points of our PCA space are indeed meaningful in reflecting the important dimensions of color term systems in the WCS.

Second, and crucially, the importance ranking of dimensions in this all-languages data set shows a very strong match to the ordering of the evolutionary stages of Fig. 1, as indicated in the caption below each chart in Fig. 4(a–e). This is the first mathematical demonstration that synchronic color naming patterns reflect the relative importance of latent color dimensions that also underlie their evolutionary emergence.

Discussion. We have shown that a straightforward application of PCA over the WCS yields dimensions of the color naming data whose maximal/minimal values correspond to focal regions of color. That is, the universal foci appear to organize the dimensions along which the data shows the most variance. Moreover, these dimensions are extracted in the order of importance of the evolutionary distinctions among color terms, confirming that naming patterns in the WCS col-

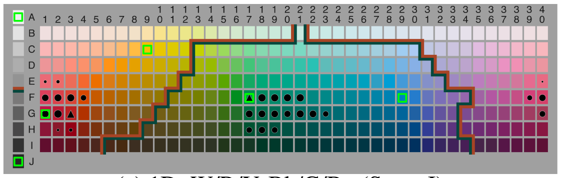

(a) 1D: W/R/Y, Bk/G/Bu (Stage I)

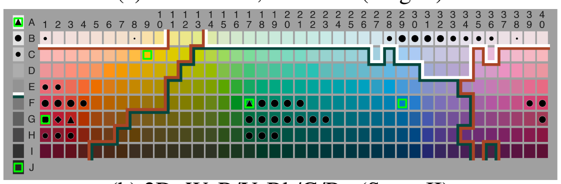

(b) 2D: W, R/Y, Bk/G/Bu (Stage II)

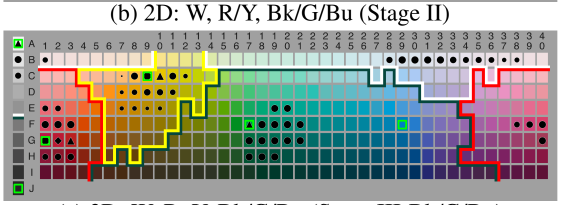

(c) 3D: W, R, Y, Bk/G/Bu (Stage III-Bk/G/Bu)

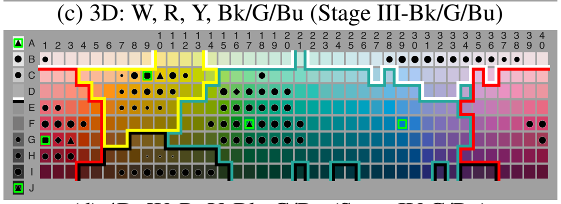

(d) 4D: W, R, Y, Bk, G/Bu (Stage IV-G/Bu)

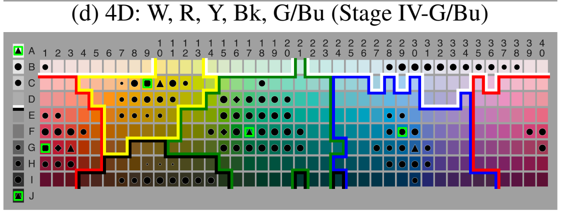

(e) 5D: W, R, Y, Bk, G, Bu (Stage V)

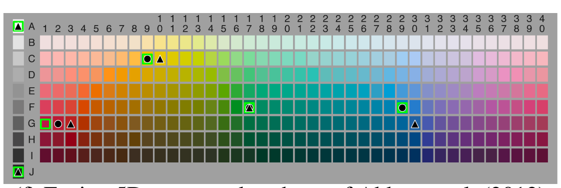

(f) Foci at 5D compared to those of Abbott et al. (2012)

Figure 4: Focal colors and associated regions from a PCA over all WCS data. (a–e) Results over each subspace for the first n dimensions; triangles = extreme points, circles = nearest neighbors, diamonds = merged extreme points, green borders = universal foci. Color labels below each chart correspond to the focal colors in each region, with the matching evolutionary stage indicated. (f) Foci in 5D, with our extreme points shown as triangles and the predicted foci of Abbott et al. (2012) as large circles.

lectively exhibit the synchronic influence of the evolutionary factors that shape the progression of color systems.

A legitimate question is whether our finding is simply the expected result of doing PCA over language data that includes languages at all the stages. Specifically, is it a simple frequency effect? That is, if languages at different stages are simply successively partitioning the data (rather than reorganizing colors in a way that changes earlier boundaries), then all languages have some boundary between warm and cool colors, all but Stage I languages have an additional boundary between white and the other warm colors, all but Stage I and II languages have another boundary between two more colors, etc. Thus, the PCA may be finding boundaries based on their frequency across languages at the different stages, rather than based on a true importance ordering.

A crucial observation that argues against this view is that the WCS contains no Stage I languages, and yet the warm–cool split of Stage I emerges as the first dimension in importance. That is: Although all languages in the WCS make *both* the warm–cool distinction of Stage I and the White–Red/Yellow distinction of Stage II, the warm–cool distinction emerges first in the PCA. This suggests that there is a detectable signal in the naming patterns that reveals the relative importance of an evolutionarily-earlier boundary over a later boundary, independently of their frequency in the data. To test this more directly, and across more stages, we next look at subsets of the languages of the WCS grouped by stage.

Analysis Over WCS Languages By Stage

We hypothesize that languages at each stage will show the same importance ranking of dimensions as found in the evolutionary progression, up to and including that stage. The set up here ensures that if we find that languages show evolutionarily-earlier distinctions as more important than later ones, this cannot be explained away as the data including languages at those earlier stages, thus skewing the frequencies toward the earlier distinctions.

Set up. In this analysis, we separately consider subsets of languages of the WCS that are in a single one of the identified evolutionary stages (Kay et al., 2009), yielding 7, 7, 41, and 14 languages at Stages II, III, IV, and V, respectively. (There are few documented Stage I languages, and none in the WCS. Also, we omit languages transitioning between stages, since they can show blends of behavior.) We perform the same PCA analysis as above, once over each of the four naming matrices limited to each stage, with the goal of seeing whether there is a match between the successive dimensions of each PCA analysis and the evolutionary stages.

Results. Tab. 1 presents the sequences of evolutionary stages revealed in the analysis of the subsets of languages by stage (omitting Stage V for space reasons). The table summarizes the color partitions in each subspace using the focal colors in each (we omit Munsell charts due to space reasons), and shows the best-matching stage from Fig. 1. (All and only dimensions accounting for $> 5\%$ of variance shown.)

Overall, the results confirm our hypothesis above: we find a very good match between the PCA analysis and the evolutionary diagram from all sets of languages, except those in Stage V. It is also the case that the majority of extreme points found in all the relevant dimensions of the four PCA analyses are at or very near focal colors. To summarize:

- All of Stages II, III, and IV show a very strong match to

Table 1: Sequences of focal colors in the Stages data.

Data set	1D subspace	2D subspace	3D subspace	4D subspace
Stage II languages	W/R/Y Bk/G/Bu ➔ **Stage I**	W R/Y Bk/G/Bu ➔ **Stage II**		
Stage III languages	W/R/Y Bk/G/Bu ➔ **Stage I**	W R/Y Bk/G/Bu ➔ **Stage II**	W R Y/warm–cool boundary G/Bu Bk ➔ **Stage III**[a,b]	W R Y G/Bu Bk warm–cool boundary ➔ **Stage III**[a,b]
Stage IV languages	W/R/Y/Bk G/Bu ➔ **Stage I**[c]	W R/Y/Bk G/Bu ➔ **Stage II**[c]	W R Y/Bk G/Bu ➔ **Stage III**[c]	W R Y G/Bu Bk ➔ **Stage IV**

[a] Stage III languages show a mix of the top and middle partitions of Stage III in 3D and 4D.

[b] Stage III languages show an additional warm/cool boundary color in 3D and 4D.

[c] Stage IV languages connect Bk with Y through the Brown region in 1D, 2D, and 3D.

the evolutionary stages.

- Stage III has, in addition to predicted extreme colors, a rainbow-like extreme area along the warm–cool boundary.
- Stage IV mostly matches the evolutionary stages, but with Yellow and Black connected through Brown in earlier dimensions. At 4D, there is a precise match to Stage IV.
- Stage V does not yield an ordering of dimensions that match the evolutionary stages. At 6D, all basic focal colors plus Purple have emerged as extreme regions.

Stage III data include languages from different sub-stages (column III in Fig. 1). Follow-up experiments with various subsets of Stage III languages reveal that the observed boundary color appears due to varying ways the different sub-stages divide up the G/Bu/Y region of color, in combination with the fact that one of the languages has an unusual basic color term for this warm–cool boundary region (Kay et al., 2009).

Stage V is a heterogeneous group, with languages having the 6 basic colors plus some number of other derived colors (13 of 14 have at least one derived color). We hypothesized that the variety in naming patterns for the non-basic colors may be swamping the signal from the basic colors. This may indicate a limitation of our method in dealing with a larger number of dimensions of color distinction. To test this, we performed the same PCA analysis over the 8 languages denoted as *approaching* Stage V (which have fewer derived colors). Here, the dimensions of the data emerged in order of the evolutionary stages, including the final stage at which Blue is distinguished from Green.

Discussion. To our knowledge, we are the first to apply a quantitative typological analysis to the languages of the WCS at the various evolutionary stages, as manually analyzed in Kay et al. (2009). Our findings provide strong support for the hypothesis that data from later stage languages can have structure that matches the evolutionary order of earlier stages. By separately analyzing languages at each specific evolutionary stage, we control for the potential frequency explanation of our results on the full WCS data set.

Further work will be required to determine the underlying causes of the cases of mismatches to the stages. Others have found, using more complex procedures, that the WCS data yield color groupings that largely, but not always, correspond to the manually derived partitionings of the color space (Lindsey & Brown, 2009; Jäger, 2012). Our method may be picking up on idiosyncratic patterns of naming, especially on smaller data sets. Regarding the Stage V data in particular, a possible shortcoming of our method is that it may not be sensitive enough to capture regularities beyond the six basic focal colors, which would be necessary to analyze this heterogeneous set of languages.

Conclusions

We present the first statistical analysis of color naming data that both shows a match between the evolutionary ordering of color systems and the importance ordering of informative dimensions of the data, and derives the focal colors from the extremes of those component dimensions. These results arise from a simple and straightforward application of PCA, a standard method from semantic typology for extracting salient dimensions from crosslinguistic naming patterns.

First, our approach reveals a quantitative importance ordering of latent dimensions of color semantics that strongly matches qualitative analyses of the evolutionary stages of color lexicons (e.g., Berlin & Kay, 1969; Kay et al., 2009). Specifically, we find that the color distinctions captured by each successive extracted dimension of the data largely correspond to the distinctions made in successive stages of color term evolution. Moreover, we show that the importance ordering of these dimensions holds even when considering languages at individual evolutionary stages, thus controlling for frequency of earlier vs. later distinctions in the data. Our work thus lends further evidence that speakers are sensitive to evolutionarily-important color distinctions that are not expressed directly by basic terms in their own language (cf. Boster, 1986; Xu et al., 2013; Gibson et al., 2017; Holmes & Regier, 2017).

Second, we find that the extreme points of the identified color dimensions correspond to a small set of focal color regions shown to occur across languages (e.g., MacLaury, 1997). Our work thus reinforces a growing body of research showing that focal colors are important dimensions of color space that serve as "anchors" for color categories (e.g., Regier et al., 2005). It has been proposed that focal colors arise at points of an irregularly-shaped perceptual space that maximize the distance between them (e.g., Jameson & D'Andrade, 1997; Regier et al., 2007). Although our method is agnostic as to the source of the latent dimensions (whether perceptual, and/or salience, as in Gibson et al., 2017, and/or communicative pressures, as in Zaslavsky et al., 2018), our results, like those of Abbott et al. (2012), show that the naming patterns of languages reflect the universal foci. Our approach further sheds light on the focal colors as extremes in the evolutionarily-important dimensions of color semantics.

Acknowledgments

JW and SS are supported by an NSERC Discovery Grant RGPIN-2017-06506 to SS. We thank the anonymous reviewers for their constructive comments.

References

Abbott, J. T., Regier, T., & Griffiths, T. L. (2012). Predicting focal colors with a rational model of representativeness. In *Proceedings of the 34th Annual Meeting of the Cognitive Science Society.*

Beekhuizen, B., Fazly, A., & Stevenson, S. (2014). Learning Meaning without Primitives: Typology Predicts Developmental Patterns. In *Proceedings of the 36th Annual Meeting of the Cognitive Science Society.*

Beekhuizen, B., & Stevenson, S. (2018). More than the eye can see: A computational model of color term acquisition and color discrimination. *Cognitive Science*, *42*(8), 2699–2734.

Berlin, B., & Kay, P. (1969). *Basic color terms: Their universality and evolution*. Berkeley, CA: UC Press.

Boster, J. (1986). Can individuals recapitulate the evolutionary development of color lexicons? *Ethnology*, *25*(1), 61–74.

Gibson, E., Futrell, R., Jara-Ettinger, J., Mahowald, K., Bergen, L., Ratnasingam, S., ... Conway, B. R. (2017). Color naming across languages reflects color use. *PNAS*, *114*(40), 10785–10790.

Heider, E. R. (1972). Universals in color naming and memory. *Journal of Experimental Psychology*, *93*(1), 10–20.

Holmes, K., & Regier, T. (2017). Categorical perception beyond the basic level: The case of warm and cool colors. *Cognitive Science*, *41*, 1135–1147.

Jäger, G. (2012). Using statistics for cross-linguistic semantics: A quantitative investigation of the typology of colour naming systems. *Journal of Semantics*, *29*(4), 521–544.

Jameson, K. A., & D'Andrade, R. G. (1997). It's not really red, green, yellow, blue: an inquiry into perceptual color space. In C. L. Hardin & L. Maffi (Eds.), *Color categories in thought and language* (pp. 295–319). CUP.

Kay, P., Berlin, B., Maffi, L., Merrifield, W. R., & Cook, R. (2009). *World color survey*. Stanford: CSLI Publications.

Levinson, S. C., Meira, S., & The Language and Cognition Group. (2003). 'Natural concepts' in the spatial topological domain – Adpositional meanings in crosslinguistic perspective: An exercise in semantic typology. *Language*, *79*(3), 485–516.

Lindsey, D. T., & Brown, A. M. (2006). Universality of color names. *PNAS*, *103*(44), 16608–16613.

Lindsey, D. T., & Brown, A. M. (2009). World color survey color naming reveals universal motifs and their within-language diversity. *PNAS*, *106*(47), 19785–19790.

MacLaury, R. E. (1997). Ethnographic evidence of unique hues and elemental colors. *BBS*, *20*(2), 202-203.

Majid, A., Boster, J. S., & Bowerman, M. (2008). The crosslinguistic categorization of everyday events: A study of cutting and breaking. *Cognition*, *109*(2), 235–250.

Regier, T., Kay, P., & Cook, R. S. (2005). Focal colors are universal after all. *PNAS*, *102*(23), 8386–8391.

Regier, T., Kay, P., & Khetarpal, N. (2007). Color naming reflects optimal partitions of color space. *PNAS*, *104*(4), 1436–1441.

Roberson, D., Davies, I., & Davidoff, J. (2000). Color categories are not universal: Replications and new evidence from a stone-age culture. *Journal of Experimental Psychology: General*, *129*, 369–398.

Xu, J., Dowman, M., & Griffiths, T. L. (2013). Cultural transmission results in convergence towards colour term universals. *Proceedings of the Royal Society B: Biological Sciences*, *280*(1758), 20123073.

Zaslavsky, N., Kemp, C., Regier, T., & Tishby, N. (2018). Efficient compression in color naming and its evolution. *PNAS*, *115*(31), 7937–7942.